

LiveTalk: Real-Time Multimodal Interactive Video Diffusion via Improved On-Policy Distillation

Ethan Chern* Zhulin Hu* Bohao Tang* Jiadi Su* Steffi Chern Zhijie Deng† Pengfei Liu†‡

SJTU SII GAIR

 Code  Models

Abstract

Real-time video generation via diffusion is essential for building general-purpose multimodal interactive AI systems. However, the simultaneous denoising of all video frames with bidirectional attention via an iterative process in diffusion models prevents real-time interaction. While existing distillation methods can make the model autoregressive and reduce sampling steps to mitigate this, they focus primarily on text-to-video generation, leaving the human-AI interaction unnatural and less efficient. This paper targets real-time interactive video diffusion conditioned on a multimodal context, including text, image, and audio, to bridge the gap. Given the observation that the leading on-policy distillation approach Self Forcing encounters challenges (visual artifacts like flickering, black frames, and quality degradation) with multimodal conditioning, we investigate an improved distillation recipe with emphasis on the quality of condition inputs as well as the initialization and schedule for the on-policy optimization. On benchmarks for multimodal-conditioned (audio, image, and text) avatar video generation including HDTF (Zhang et al., 2021), AVSpeech (Ephrat et al., 2018), and CelebV-HQ (Zhu et al., 2022), our distilled model matches the visual quality of the full-step, bidirectional baselines of similar or larger size with 20× less inference cost and latency. Further, we integrate our model with audio language models and long-form video inference technique Anchor-Heavy Identity Sinks to build LiveTalk, a real-time multimodal interactive avatar system. System-level evaluation on our curated multi-turn interaction benchmark shows LiveTalk outperforms state-of-the-art models (Sora2, Veo3) in multi-turn video coherence and content quality, while reducing response latency from 1 to 2 minutes to real-time generation, enabling seamless human-AI multimodal interaction.

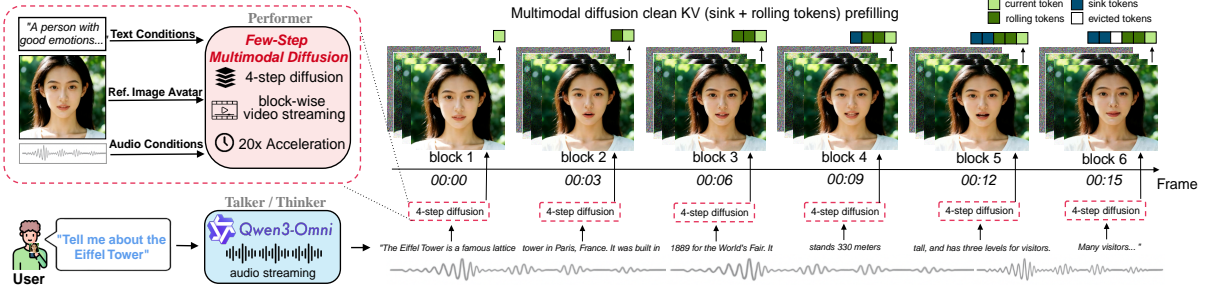


Figure 1: **Overview of the LiveTalk system.** Given a user audio/text query, Qwen3-Omni (Xu et al., 2025) processes the query and generates streaming audio responses in real-time. Our few-step multimodal diffusion model takes the streaming audio along with the reference image avatar and text conditions to generate synchronized video responses through block-wise AR generation. Each block (3 latent frames) undergoes 4-step diffusion, achieving 20× acceleration over the baseline (See Tab. 1, Ours vs. OmniAvatar-1.3B (Gan et al., 2025)). To support long-horizon streaming with sub-second latency, we perform clean KV prefilling across blocks using a *sink+rolling* token cache: persistent sink tokens retain global context, rolling tokens carry recent history.

* Equal contribution.

† Co-advising.

‡ Corresponding author.

1 Introduction

Diffusion transformers (DiTs) (Peebles and Xie, 2023; Brooks et al., 2024; Wan et al., 2025; Chen et al., 2025) have enabled appealing visual fidelity for video generation. The sampling involves denoising all video frames simultaneously using bidirectional attention via an iterative process. It can be costly, e.g., state-of-the-art models such as Veo3 (Google DeepMind, 2025) and Sora2 (OpenAI, 2025) require 1 to 2 minutes to generate a single 5 to 10 second clip, incurring prohibitive inference expense and creating fundamental barriers to real-time applications.

As a result, there is increasing interest in converting pretrained, bidirectional, many-step video diffusion models into causal, few-step autoregressive (AR) ones via distillation techniques (Lu and Lab, 2025; Agarwal et al., 2024; Yin et al., 2024b; Huang et al., 2025; Liu et al., 2025). The resultant models enable low-latency streaming generation. However, systematic investigation of complex multimodal conditioning for distillation remains largely underexplored, particularly the simultaneous integration of text, images, and audio for interactive avatar generation. Such natural conditioning is critical for building general-purpose interactive AI systems that can not only *understand* across modalities but also *express themselves visually* for natural human-AI interaction.

This paper aims to establish a real-time interactive video diffusion model conditioned on a multimodal context. We first perform an in-depth investigation based on the leading on-policy distillation approach, Self Forcing (Huang et al., 2025), which involves an initialization stage based on ODE trajectory distillation (Song et al., 2023; Yin et al., 2025; Gu et al., 2023; Berthelot et al., 2023) to obtain a few-step causal student model with block-wise causal attentions, as well as a distribution matching distillation (DMD) (Yin et al., 2024b,a, 2025) stage to minimize *exposure bias* (Ning et al., 2023; Schmidt, 2019; Huang et al., 2025) from causality based on on-policy rollouts.

We observe Self Forcing can result in extensive visual artifacts, e.g., flickering effects (see row 1 of Fig. 4) for the multimodal setting, with related issues reported in the community (Chen, 2025).

We speculate that the issues may stem from the complex interplay among the inherent multiple components in DMD, which renders the optimization unstable and fragile, particularly under complex multimodal conditions.

To address this, we investigate an improved distillation recipe with emphasis on the quality of the conditions as well as the initialization and schedule for a stable on-policy optimization. Concretely, we advocate (1) **refining multimodal conditions for distillation**, e.g., making the condition image high-quality and the text prompts motion-focused; (2) **training ODE initialization to convergence** before applying on-policy DMD training; and (3) **maximizing learning within DMD’s limited learning window** through aggressive learning rates and tuned classifier-guidance (CFG) (Ho and Salimans, 2022) scales. We validate these by distilling a multimodal variant of Wan2.1 (Wan et al., 2025) (i.e., OmniAvatar (Gan et al., 2025)) and benchmarking on diverse multimodal-driven avatar generation datasets, including HDTF (Zhang et al., 2021), AVSpeech (Ephrat et al., 2018), and CelebV-HQ (Zhu et al., 2022).

Our distilled model even surpasses some 5B and 14B bidirectional, many-step baselines in multiple aspects. The distillation process dramatically improves inference efficiency: achieving 24.82 FPS ($20\times$ speedup compared to the vanilla model), and reducing first-frame latency to subsecond ($200\times$ speedup), opening the door for real-time interactive communication.

Based on the distilled model, we build a real-time multimodal avatar system, LiveTalk, that enables seamless interaction between humans and AI. The system (Fig. 1) leverages existing audio language models (Xu et al., 2025) for reasoning and speech, while our model renders talking avatars in real-time with high visual fidelity (Fig. 1). To preserve speaker identity in long-horizon video streaming, we introduce a training-free technique, the Anchor-Heavy Identity Sinks (AHIS), which successfully keeps the generated speaker visually undistorted on a time scale of minutes. We also curate a new multi-turn interaction benchmark for this new form of real-time multimodal talking avatar systems. Evaluations against state-of-the-art models Veo3 (Google DeepMind, 2025) and Sora2 (OpenAI, 2025) show that our system substantially outperforms them in multi-round coherence and content quality, while reducing response latency from minutes to real-time generation, enabling truly interactive communication. In summary, our contributions are:

- **Actionable distillation framework for multimodal video diffusion.** We establish a systematic recipe for training real-time multimodal interactive video models conditioned on text, image, and audio. We identify three key improvements for stable on-policy distillation under complex multimodal conditions: curated multimodal conditioning, converged ODE initialization, and aggressive optimization schedule. Our ablation study (Tab. 3) demonstrates these collectively deliver significant quality improvements across perceptual metrics, audio-visual synchronization, and aesthetic quality.
- **Real-time multimodal video generation with $20\times$ speedup.** Our distilled 4-step model matches or exceeds bidirectional diffusion baselines while reducing inference cost by over $20\times$. Our 1.3B model matches or surpasses the 1.3B bidirectional variant (OmniAvatar-1.3B (Gan et al., 2025)) and larger baselines, including Hallo3 (Cui et al., 2024), FantasyTalking (Wang et al., 2025), and AniPortrait (Wei et al., 2024) across quality metrics, while achieving real-time generation at 24.82 FPS on a single GPU.

- **Complete real-time multimodal interactive avatar system.** We build LiveTalk and propose a benchmark for evaluating multi-turn interaction quality of multimodal interactive avatar system. Our system outperforms Veo3 (Google DeepMind, 2025) and Sora2 (OpenAI, 2025) in multi-video coherence and content quality metrics, while achieving sub-second response, enabling seamless human-AI interaction.

2 Related Work

2.1 Multimodal Video Diffusion

Modern video synthesis has evolved beyond text control to incorporate richer conditioning signals such as images and audio, enhancing controllability and versatility (Wan et al., 2025; Kong et al., 2024; Chen et al., 2025; Gan et al., 2025). These multimodal capabilities enable applications including image-guided video editing (Wan et al., 2025; Kong et al., 2024) and multimodal-driven virtual avatars (Gan et al., 2025; Chen et al., 2025). Extending video diffusion to handle complex multimodal conditions presents challenges including architectural designs (Ju et al., 2025; Lin et al., 2025; Hu et al., 2025), training overheads (Wan et al., 2025; Kong et al., 2024), and cross-modal alignment (Li et al., 2024; Gan et al., 2025; Chen et al., 2025). However, existing multimodal video diffusion predominantly adopts pure diffusion paradigms requiring many-step iterative denoising across entire sequences. While achieving high visual quality, their significant inference cost and latency render them impractical for real-time applications. While recent work (Low and Wang, 2025) has begun exploring real-time audio-driven models, comprehensive recipes for training stability, exposure bias mitigation, and rigorous benchmarking remain largely underexplored.

2.2 Real-Time Video Diffusion

Real-time video diffusion adopts hybrid modeling combining AR and diffusion components: AR enables streaming generation without future frame dependencies, while diffusion ensures high visual fidelity (Teng et al., 2025; Bruce et al., 2024; Liu et al., 2025; Yang et al., 2025). However, AR video diffusion faces *exposure bias* (Ning et al., 2023; Schmidt, 2019; Huang et al., 2025) from error accumulation during autoregression. Post-training strategies address this through distilling guidance from bidirectional models (Yin et al., 2025), increasing robustness to imperfect frames (Chen et al., 2024), and mitigating train-test gaps via self-generated rollouts (Huang et al., 2025). However, these techniques have primarily been explored for text-to-video and remain largely unexamined for multimodal video diffusion. We bridge this gap by systematically investigating and improving on-policy distillation for multimodal conditions (text, image, audio), emphasizing condition quality, initialization, and optimization dynamics, and establishing a comprehensive multi-turn interaction benchmark to guide future development for real-time avatar systems.

3 Improved Distillation for Real-Time Multimodal Interactive Video Diffusion

This section briefly reviews Self Forcing (Huang et al., 2025), a dominant distillation method for constructing real-time video diffusion, discusses its limitations for handling video diffusion models with multimodal conditioning, and elaborates on our improved distillation strategies.

3.1 Preliminary: Self Forcing

There has been ongoing interest in transferring bidirectional, many-step, pure diffusion models (Wan et al., 2025; Chen et al., 2025) into causal, few-step AR ones via distillation-based post-training approaches (Chen et al., 2024; Huang et al., 2025; Yin et al., 2025).

Self Forcing (Huang et al., 2025) is one of the most effective approaches, which follows a two-stage procedure. First, an ODE initialization (Song et al., 2023; Yin et al., 2025; Gu et al., 2023; Berthelot et al., 2023) stage is performed to obtain a few-step causal student model with block-wise causal attentions. Then, it conducts on-policy distillation (Lu and Lab, 2025; Agarwal et al., 2024) with self-generated rollouts (Huang et al., 2025) using the distribution matching distillation (DMD) principle (Yin et al., 2024b,a, 2025), which helps minimize the *exposure bias* (Ning et al., 2023; Schmidt, 2019; Huang et al., 2025) of the student model.

Concretely, let \mathbf{x}_0 denote the latent corresponding to video frames yielded by a compression variational auto-encoder (VAE) (Kingma and Welling, 2013; Rombach et al., 2022). Letting \mathbf{c} denote the conditioning for generation. Self Forcing aims at distilling a vanilla teacher model into a student one g_ϕ that can generate in a block-by-block manner, where each block consists of multiple latent frames (e.g., 3 in our case) in real-time.

ODE Initialization first performs trajectory distillation. It subsamples $k = 4$ timesteps from the teacher’s $N = 48$ step denoising trajectory $\{\mathbf{x}_{t_j}\}_{j=0}^N$, based on which the causal student is asked to predict clean \mathbf{x}_0 :

$$\mathcal{L}_{\text{ODE}} = \mathbb{E}_{t \sim \{t_{iN/k}\}_{i=0}^{k-1}} \left[\sum_b \|g_\phi(\mathbf{x}_t^b, t, \mathbf{c}) - \mathbf{x}_0^b\|_2^2 \right], \quad (1)$$

with the superscript b denoting the b -th block.

Distribution Matching Distillation (DMD) then addresses the exposure bias (Yin et al., 2024b,a, 2025) of the model after ODE initialization due to the training under teacher-forced trajectories. The algorithmic implementation introduces two extra models: a frozen teacher score network s_θ (Song et al., 2020; Yin et al., 2024b,a) and a trainable critic s_ψ , and the training alternates between updating the generator g_ϕ and the critic s_ψ . Specifically, the gradient for updating the student model is

$$-\mathbb{E}_{\tau, \hat{\mathbf{x}}_0, \mathbf{x}_\tau} \left[\left(s_\theta(\mathbf{x}_\tau, \tau, \mathbf{c}) - s_\psi(\mathbf{x}_\tau, \tau, \mathbf{c}) \right) \frac{\partial \hat{\mathbf{x}}_0}{\partial \phi} \right], \quad (2)$$

where $\hat{\mathbf{x}}_0 = g_\phi(\mathbf{z}, 0, \mathbf{c})$ with $\mathbf{z} \sim \mathcal{N}(0, I)$ and \mathbf{x}_τ comes from adding τ -time noise to $\hat{\mathbf{x}}_0$. Note that, for multimodal conditioning, s_θ can be estimated with the classifier-free guidance (Ho and Salimans, 2022) strategy using separate scales for various conditions. The critic s_ψ learns to tracking the evolving distribution of the generator by minimizing the standard diffusion denoising objective:

$$\mathcal{L}_{\text{critic}} = \mathbb{E}_\tau \left[\|s_\psi(\mathbf{x}_\tau, \tau, \mathbf{c}) - \hat{\mathbf{x}}_0\|_2^2 \right]. \quad (3)$$

3.2 Issues on Existing Distillation Recipe

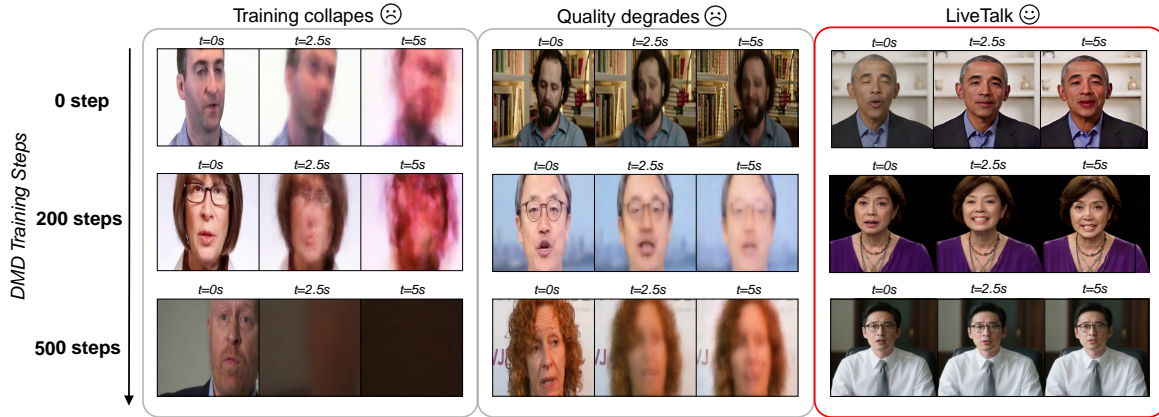


Figure 2: **Degraded training performance with Self Forcing DMD.** Left and middle columns show failure cases from Self Forcing DMD training, exhibiting quality degradation. Right column shows stable results from our method.

When naively applying on-policy distillation with default settings from Self Forcing (Huang et al., 2025) to distill multimodal video diffusion models, we encounter significant training instabilities that manifest as visual artifacts (Fig. 2, left and middle).

This challenge likely stems from the complex interplay in the DMD training, where the critic score network s_ψ learns to denoise the noised generator rollout \mathbf{x}_τ in the critic training stage, creating a delicate interdependency. When the generator output degrades catastrophically, the critic score network s_ψ receives corrupted training signals, resulting in inaccurate gradient estimates that further degrade the generator, potentially triggering mode collapse (Ge et al., 2025). While Self Forcing demonstrates robustness under text conditioning, multimodal conditions introduce additional complexity that could amplify instability (Chen, 2025). Through investigative studies, we motivate three critical factors that contribute to the instability:

Data Quality Issues. In our initial trials, we selected 2000 multimodal conditions (reference image and audio) from each of the Hallo3 and HDTF datasets and applied distillation with default settings of Self Forcing. However, the DMD training collapsed after several hundred iterations, with outputs degrading to black images (Fig. 2, left). Through investigation, we discovered that the quality of the reference image condition critically influences distillation stability. Specifically, existing datasets often contain lower-quality images with artifacts (such as Hallo3’s overall low image quality and HDTF’s facial blurriness) which lead to imperfect generator rollouts during distillation. These imperfect rollouts in turn produce corrupted training signals that destabilize the learning process. To address this issue, we filtered distillation training samples using brightness and quality metrics, which yielded slight improvements in training stability (Fig. 2, middle). However, we still observed blurriness emerging after a few hundred training steps, suggesting that additional factors contribute to distillation instability, as we discuss below.

Insufficient ODE Initialization. We observe that insufficient ODE initialization creates a weak foundation that leads to instability during subsequent on-policy DMD training. Unlike text-to-video distillation (Huang et al.,

2025) where even ODE checkpoints generating low-quality videos (User, 2024) can lead to successful distillation, multimodal-conditioned distillation exhibits markedly different behavior. Insufficient ODE training manifests as severity-dependent DMD failures: collapse in extreme cases, and performance plateaus with blurry artifacts in moderate cases. This suggests that distilling multimodal video diffusion requires a more robust starting point to facilitate the critic-generator interdependency to function stably.

Limited Learning Window. We observe that the effective learning window for multimodal-conditioned DMD training is considerably short, with the model reaching peak performance within a few hundred steps before degrading. In contrast to text-to-video distillation where prior work reports convergence in 90 minutes, multimodal conditioning exhibits a peak-then-degrade pattern (Chen, 2025). With standard learning rates and guidance scales, the model fails to sufficiently learn optimal multimodal alignment (particularly audio-visual synchronization) before degradation occurs, leaving considerable performance on the table.

These three factors interact to create a relatively fragile on-policy distillation training procedure that careful treatment beyond what is required for text-to-video distillation.

3.3 Improvements on Existing Distillation Recipe

To address the challenges identified in Section 3.2, we propose three key improvements to the on-policy distillation procedure for multimodal diffusion models. We demonstrate the effectiveness of each component through ablations in Tab. 3 and Fig. 4.

Refining Multimodal Conditions for Distillation. Instead of using existing dataset directly, we meticulously curate high-quality multimodal conditions $\mathbf{c} = \{\mathbf{c}_{\text{text}}, \mathbf{c}_{\text{img}}, \mathbf{c}_{\text{audio}}\}$ to provide clean training signals. Note that we mainly consider improving \mathbf{c}_{img} and \mathbf{c}_{text} in this paper. We adopt targeted curation strategies for different dataset. For Hallo3, characterized by overall low image quality, we employ Qwen-Image (Wu et al., 2025) to generate semantically consistent yet high-quality reference frames \mathbf{c}_{img} . For HDTF, which primarily suffers from facial blurriness, we apply super-resolution (Zhou et al., 2022) to obtain clear facial details. Furthermore, we utilize Qwen2.5-VL-72B (Bai et al., 2025b) to refine text prompts \mathbf{c}_{text} , emphasizing dynamic motion and facial expressions to enrich temporal and semantic information. **Training ODE Initialization to Convergence.** To establish a robust starting point for on-policy distillation, we train ODE initialization to full convergence using an extended training schedule, ensuring that the student model has thoroughly learned to denoise across all timesteps before transitioning to DMD. This convergent ODE checkpoint provides a strong foundation that stabilizes the delicate critic-generator interdependency during subsequent on-policy training.

Maximizing Learning Within the Limited Window of On-Policy Distillation. To maximize learning before degradation occurs, we employ an aggressive learning rate schedule ($2\times$ baseline) that accelerates convergence within the limited effective learning window. We also apply higher CFG guidance for the teacher model to significantly strengthen audio conditioning for lip synchronization. While these strategies introduce potential instability, they represent a necessary trade-off to achieve optimal multimodal conditioning before the peak-then-degrade transition. These aggressive training strategies enable the model to learn strong audio-visual alignment (evaluated by Sync-C and Sync-D) while maintaining high visual quality.

4 Building Real-Time Multimodal Interactive Systems

Based on the distilled model, we build LiveTalk (Fig. 1), a complete real-time multimodal interactive avatar system integrating our model with Qwen3-Omni (Xu et al., 2025) for end-to-end visual communication. Our modular architecture comprises two key components: a real-time performer module (our distilled video diffusion model) that renders synchronized talking avatars, and a thinker/talker module (Qwen3-Omni) that handles reasoning and generates streaming audio responses. We detail the system pipeline overview, and the two critical system-level challenges below.

System Pipeline Overview. Figure 1 illustrates LiveTalk’s interaction pipeline. When a user provides audio or text input, Qwen3-Omni (Xu et al., 2025) generates streaming audio responses. Our video diffusion model takes three multimodal conditions: (1) streaming audio output $\mathbf{c}_{\text{audio}}$ from Qwen3-Omni, (2) reference image avatar \mathbf{c}_{img} defining visual identity, and (3) text prompts \mathbf{c}_{text} describing the desired motion (e.g., “A person speaking naturally with expressive gestures and emotions”). These jointly drive block-wise autoregressive generation, where each block of $b = 3$ latent frames undergoes $k = 4$ diffusion steps. Clean KV cache from previous blocks is prefilled for visual consistency, enabling streaming synchronized to audio with sub-second first-frame latency.

Streaming Audio Conditioning. Audio-driven video generation requires acoustic context from adjacent frames for smooth lip-sync and natural motion. Conditioning each video block solely on temporally aligned audio causes discontinuities at block boundaries. Waiting for extended audio sequences (encoding entire clips before generation) introduces prohibitive latency. Our solution uses overlapped windowing: we encode and generate as soon as a small windowed segment becomes available, providing rich acoustic context while maintaining real-time responsiveness.

Long-Horizon Video Streaming with Speaker Identity Preservation. Interactive dialogues often last for minutes, which requires the generated avatar’s identity in the generated video to remain consistent over long time spans. Although Self Forcing (Huang et al., 2025) reduces the train–test mismatch by aligning the distributions of training and inference, identity still degrades once generation extends beyond the training window: accumulated errors can cause color drift and geometric distortions. This motivates a simple question: within a fixed attention window, can we downweight error-prone recent blocks and upweight high-fidelity identity representations instead?

Following this intuition, we propose a training-free method, **Anchor-Heavy Identity Sinks (AHIS)**: we allocate part of the KV cache as attention sinks (Xiao et al., 2024) that permanently store early high-fidelity speaker frames as identity anchors, while the remaining rolling KV tokens maintain contextual continuity and prevent visual discontinuities. Unlike standard attention sink designs, AHIS deliberately allocates a much larger fraction of the KV window to identity sink tokens than to rolling tokens. Under the same KV-cache budget, increasing the proportion of sink tokens while reducing rolling KV tokens simultaneously strengthens the focus on high-fidelity identity and suppresses attention to accumulated errors in subsequent generations. Specifically, we set the KV window to 5 blocks, with the first three blocks used as sink tokens and the last two as rolling KV tokens. In our experiments, this setting effectively preserved the speaker’s appearance over several minutes of generated video.

Parallel Pipeline for Video Denoising and Decoding. Streaming video generation requires diffusion denoising (predicting clean latents) and VAE decoding (converting to pixels). Sequential execution risks playback stalling when generation time exceeds playback duration. We adopt pipeline parallelism: while the current block undergoes denoising, the previous block is simultaneously decoded. This reduces per-block latency from the sum to the maximum of the two stages, ensuring generation stays ahead of playback and enabling non-stalling streaming with seamless real-time rendering.

5 Experiments

5.1 Settings

Model and Training Configuration. We conduct distillation experiments on the multimodal variant of Wan2.1 (Wan et al., 2025), instantiated by OmniAvatar (Gan et al., 2025). OmniAvatar is a multimodal bidirectional diffusion model that accepts text, image, and audio as conditional inputs to synthesize videos with realistic facial expressions and lip synchronization.

We adapt OmniAvatar-1.3B into a causal student model by modifying its architecture to support causal attention and KV cache for AR generation. Following the block-wise AR paradigm, we adopt a block size of $b = 3$ latent frames, enabling efficient streaming generation while maintaining temporal coherence.

For training data, we curate 4000 multimodal conditions from Hallo3 (Cui et al., 2024) and the HDTF training split (Zhang et al., 2021). Note that we directly use the audio samples from the dataset, and refine the image and text conditions following the approach in Section 3.3. Using the bidirectional OmniAvatar-1.3B, we generate corresponding ODE trajectories from these curated conditions for ODE initialization training.

We apply the combined distillation strategy described in Section 3.3 for on-policy DMD training. Specifically, we use OmniAvatar-14B as the frozen teacher score network s_θ , and OmniAvatar-1.3B as the trainable critic model s_ψ . The generator is initialized via ODE trajectory matching, trained for 20k steps on the generated trajectory pairs. Subsequently, DMD distillation is performed using the same 4000 curated conditions, reaching optimal generation quality within approximately 1000 steps. Detailed hyperparameters and training configurations are provided in the Appendix.

System Integration. To build a complete real-time multimodal interactive avatar system, we integrate our distilled video generation model with Qwen3-Omni (Xu et al., 2025), which serves as the “thinker/talker” module responsible for reasoning and generating streaming audio responses. During interaction, the system takes the streaming audio output from Qwen3-Omni (Xu et al., 2025), along with a reference image avatar and text prompt, as conditioning inputs to our video diffusion model for real-time video generation (Fig. 1).

Evaluation Protocol. We conduct evaluations under two settings: (1) **single-round evaluation**: standard multimodal-driven avatar video generation with a single reference image and audio clip, evaluated on established benchmarks (HDTF, AVSpeech, CelebV-HQ); and (2) **Multi-round evaluation**: free-form conversational interaction between users and the multimodal avatar system, assessed using our proposed multi-turn interaction benchmark spanning 9 evaluation dimensions.

5.2 Single-Round Evaluation

We evaluate our distilled model against several baselines including AniPortrait (Wei et al., 2024), Hallo3 (Cui et al., 2024), FantasyTalking (Wang et al., 2025), and the bidirectional teachers OmniAvatar-1.3B and OmniAvatar-14B (Gan et al., 2025). The evaluation is conducted on 100 randomly sampled 5-second clips across three multimodal-driven avatar video generation benchmarks: HDTF test split (Zhang et al., 2021) for in-domain

Table 1: Quantitative comparison with existing multimodal avatar generation methods on the test set. Our distilled model achieves comparable or *superior* visual quality, aesthetics, and lip-sync performance to the bidirectional baseline OmniAvatar-1.3B, while delivering approximately $20\times$ throughput speedup and $200\times$ faster first-frame latency. **Bold** indicates best performance.

Methods	Model Size	Throughput (FPS) \uparrow	Latency (s) \downarrow	FID \downarrow	FVD \downarrow	Sync-C \uparrow	Sync-D \downarrow	IQA \uparrow	ASE \uparrow
HDTF									
Ground Truth	-	-	-	-	-	8.06	7.31	4.20	2.56
AniPortrait (Wei et al., 2024)	2.5B	0.38	211.77	21.41	323.08	1.17	13.32	4.03	2.40
Hallo3 (Cui et al., 2024)	5B	0.21	589.69	23.18	290.05	6.01	9.34	3.54	2.02
FantasyTalking (Wang et al., 2025)	14B	0.35	232.06	24.54	488.53	2.92	12.21	3.78	2.14
OmniAvatar-14B (Gan et al., 2025)	14B	0.20	412.50	8.82	151.30	6.29	9.16	3.91	2.30
OmniAvatar-1.3B (Gan et al., 2025)	1.3B	0.97	83.44	10.85	187.46	3.85	11.38	3.87	2.25
Ours	1.3B	24.82	0.33	13.68	190.07	4.50	11.00	4.13	2.52
AVSpeech									
Ground Truth	-	-	-	-	-	5.83	8.19	4.11	2.50
AniPortrait (Wei et al., 2024)	2.5B	0.38	211.77	32.14	553.27	0.84	12.86	4.07	2.49
Hallo3 (Cui et al., 2024)	5B	0.21	589.69	32.88	502.05	5.10	9.50	3.49	1.99
FantasyTalking (Wang et al., 2025)	14B	0.35	232.06	33.08	509.56	2.16	11.72	3.68	2.15
OmniAvatar-14B (Gan et al., 2025)	14B	0.20	412.50	31.21	450.09	5.18	9.17	3.77	2.25
OmniAvatar-1.3B (Gan et al., 2025)	1.3B	0.97	83.44	31.55	483.94	3.16	10.91	3.70	2.18
Ours	1.3B	24.82	0.33	33.96	486.77	3.71	10.79	4.08	2.50
CelebV-HQ									
Ground Truth	-	-	-	-	-	5.36	8.13	4.31	2.82
AniPortrait (Wei et al., 2024)	2.5B	0.38	211.77	23.87	427.43	1.14	12.07	4.24	2.79
Hallo3 (Cui et al., 2024)	5B	0.21	589.69	27.15	432.18	4.65	9.20	3.77	2.35
FantasyTalking (Wang et al., 2025)	14B	0.35	232.06	26.87	454.98	2.48	10.95	3.93	2.46
OmniAvatar-14B (Gan et al., 2025)	14B	0.20	412.50	21.25	406.78	4.81	8.85	4.03	2.59
OmniAvatar-1.3B (Gan et al., 2025)	1.3B	0.97	83.44	22.28	382.35	3.09	10.40	3.98	2.52
Ours	1.3B	24.82	0.33	25.37	437.97	3.78	10.08	4.29	2.79

evaluation, and AVSpeech (Ephrat et al., 2018) and CelebV-HQ (Zhu et al., 2022) for out-of-domain evaluation. Following the evaluation protocol in (Gan et al., 2025; Chen et al., 2025), we employ FID (Heusel et al., 2018), FVD (Unterthiner et al., 2019), IQA (Wu et al., 2023), and ASE (Wu et al., 2023) to assess visual quality and aesthetics, and Sync-C/D (Chung and Zisserman, 2017) to measure lip-sync synchronization between the audio condition and the generated lip movements. Ground truth videos are also evaluated on reference-free metrics for comparison. To assess inference efficiency, we measure the throughput and first-frame latency of all models at 512×512 resolution on a single GPU. The evaluation protocol consists of several warm-up generations followed by multiple test runs with randomly sampled conditions, with all metrics averaged across the test runs to ensure statistical reliability.

Results and Analysis. Our distilled model achieves comparable or *superior* visual quality, aesthetics, and lip-sync synchronization compared to the bidirectional many-step variant (i.e., OmniAvatar-1.3B (Gan et al., 2025)) across in-domain HDTF (Zhang et al., 2021) and out-of-domain AVSpeech (Ephrat et al., 2018) and CelebV-HQ benchmarks. More significantly, our model demonstrates substantial efficiency gains: 24.82 FPS throughput compared to 0.97 FPS ($25\times$ speedup) and first-frame latency reduced from 83.44s to 0.33s ($250\times$ faster). Remarkably, our 1.3B distilled model achieves comparable or superior performance to several larger bidirectional, many-step models, including AniPortrait (2.5B), Hallo3 (5B), and FantasyTalking (14B), while maintaining significantly higher efficiency (over $100\times$ improvement in latency and $50\times$ in throughput). Figure 3 shows representative video samples generated by our model.

5.3 Multi-Round Evaluation

Existing evaluation benchmarks for multimodal-driven avatar video generation primarily focus on single-audio-clip metrics, using traditional visual evaluation measures such as lip-sync accuracy (Sync-C/D), image quality (FID/FVD/IQA), and aesthetics (ASE). However, these metrics fail to assess multi-turn interaction quality, a critical requirement for real-world multimodal conversational applications. To address this gap, we propose a multi-round interaction benchmark and an evaluation protocol based on vision-language models (VLMs) (Bai et al., 2023; Lee et al., 2024) to comprehensively evaluate the multi-turn conversational capabilities of audio-driven avatar systems. **Benchmark Design Methodology.** We meticulously curate 100 multi-turn evaluation scenarios that require multimodal interactive AI systems to provide coherent audio-driven video generation responses across conversational turns. For example, in the first round, a user might ask, “Tell me about the Eiffel Tower,” and the system should generate a coherent introduction to the landmark with synchronized audio and video. In the subsequent round, the user might follow up with, “Where is this avenue located?” referencing the Champs-Élysées mentioned in the previous response. The AI system must generate coherent video and audio outputs that are contextually grounded in its previous multimodal conversational history, demonstrating the ability to maintain temporal coherence and contextual awareness across multiple interaction turns.

Evaluation Protocol. We adopt Qwen3-VL-30B-A3B-Instruct (Bai et al., 2025b) as the VLM-as-evaluator (Lee et al., 2024) with structured prompts tailored to each dimension. Our evaluation framework comprises two primary categories: *Visual Interaction Performance* (4 dimensions: Emotional Appropriateness, Nonverbal Interaction,




Figure 3: **Examples of multimodal-conditioned avatar video generation by our model.** Our model generates temporally coherent video with natural facial expressions, accurate lip-sync to the audio conditions, and consistent visual identity across frames.

Multi-Video Coherence, and Conversational Naturalness) and *Interaction Content Quality* (5 dimensions: Semantic Relevance, Information Completeness, Logical Consistency, Context Understanding, and Overall Interaction Experience). Detailed evaluation dimension descriptions are provided in the Appendix. System-generated audio responses are transcribed via FunASR (Gao et al., 2023) to enable audio-video synchronization and content quality evaluation.

To enable fair comparison across baselines, we normalize raw scores for each metric using z-score transformation. Specifically, for each metric, we create a single pooled distribution containing scores from all methods on all evaluation samples, compute z-scores as $z = (x - \mu)/\sigma$ where μ and σ are from the pooled distribution, and convert each z-score to its percentile (0-100). We report average percentiles per baseline for each metric, providing relative performance interpretation.

Table 2: Multi-Round Interaction Quality Evaluation. Our method is benchmarked against baselines on the proposed interaction benchmark. We report Z-Score percentile values for Visual Interaction Performance and Interaction Content Quality metrics, along with average inference throughput and latency per turn. **Bold** indicates best performance, and underline indicates second-best.

Method	Throughput (FPS) \uparrow	Latency (s) \downarrow	Visual Interaction Performance \uparrow				Interaction Content Quality \uparrow				
			EA	NI	MVC	CN	SR	IC	LC	CU	OIE
Veo3	-	61.46	23.51	24.68	26.68	24.93	25.17	17.21	20.56	20.05	19.07
Sora2	-	121.85	75.78	72.20	25.85	74.32	66.68	50.01	65.07	68.80	53.09
LiveTalk	24.82	1.16	<u>59.54</u>	<u>60.42</u>	87.26	<u>56.32</u>	72.02	81.27	75.20	72.65	81.59

Note: EA: Emotional Appropriateness, NI: Nonverbal Interaction, MVC: Multi-Video Coherence, CN: Conversational Naturalness; SR: Semantic Relevance, IC: Information Completeness, LC: Logical Consistency, CU: Context Understanding, OIE: Overall Interaction Experience.

Results and Analysis. Table 2 presents multi-round evaluation against Veo3 (Google DeepMind, 2025) and Sora2 (OpenAI, 2025). LiveTalk outperforms both baselines in multi-video coherence and content quality metrics while remaining competitive on other visual interaction dimensions, demonstrating its advantages in maintaining meaningful and coherent multi-turn multimodal interaction. Notably, both Veo3 and Sora2 exhibit visual drift across turns due to prolonged generation times (61-122s) that break conversational flow and lack of effective memory mechanisms for multi-round video interaction. In contrast, LiveTalk adopts AR generation with KV cache to maintain visual memory states, while also leveraging the Qwen3-Omni thinker/talker module to preserve textual memory states, enabling coherent generation across modalities while preserving real-time responsiveness.

Table 3: Ablation study showing the impact of various improvements. Each row sequentially adds one component: (1) curated high-quality multimodal conditions, (2) training ODE distillation to convergence, (3) accelerated learning rate, and (4) tuned CFG scale for real-score estimation. "Final Configuration" is the best performing configuration after all improvements are applied. The last row represents the baseline multimodal conditions with all other improvements applied.

Setting	FID↓	FVD↓	Sync-C↑	Sync-D↓	IQA↑	ASE
Baseline	27.10	338.08	3.13	11.77	3.95	2.38
+ Curated Multimodal Conditions	14.90	217.68	3.53	11.47	3.99	2.31
+ Converged ODE Initialization	11.67	169.75	4.15	11.19	4.18	2.56
+ Aggressive Learning Rate	12.10	179.73	4.29	11.07	4.15	2.53
+ Tuned Teacher Score CFG (Final Configuration)	13.68	190.07	4.50	11.00	4.13	2.52
Final Configuration without Curated Multimodal Conditions	23.89	261.19	3.85	11.42	4.06	2.47

5.4 Ablations

We conduct ablation studies on the four proposed components in Section 3.3: (1) refined multimodal conditions, (2) converged ODE initialization, (3) aggressive learning rate schedule, and (4) tuned teacher score CFG guidance. We evaluate on the HDTF test set using 100 randomly sampled 5-second clips, measuring both visual quality metrics (FID, FVD, IQA, ASE) and audio-visual synchronization metrics (Sync-C, Sync-D). Table 3 presents the results, with each row sequentially adding one component to demonstrate its incremental contribution.




Figure 4: Ablation study visualization. Generated video shows progressive improvements for each ablated component: (1) curated multimodal conditions, (2) ODE initialization for full convergence (20k steps), and (3) aggressive hyperparameter settings (doubled learning rates, CFG=6).

Results and Analysis. We systematically evaluate the impact of each design choice on distillation quality.

Baseline. Our baseline uses a combined of 4000 multimodal conditions (2000 from Hallo3, and 2000 from HDTF), selected based on brightness and image quality metrics of the reference image. We then follow default Self Forcing distillation settings: ODE initialization for 4000 steps, followed by 1000 DMD steps (critic learning rate 4×10^{-7} , generator learning rate 2×10^{-6} , teacher score CFG scale 4) with EMA decay 0.99. This configuration has the poorest performance, with generated videos exhibiting severe quality degradation.

Curated Multimodal Conditions. Training on curated multimodal conditions, the distilled model generates videos with substantially improved visual quality and audio-visual synchronization. While severe degradation is eliminated, inter-frame artifacts including blurriness and flickering remain noticeable.

Converged ODE Initialization. Extending ODE initialization to 20000 steps for full convergence, the model completely resolves these visual defects, with all quality metrics reaching peak performance at this stage.

Aggressive LR & Tuned Teacher Score CFG. Finally, we explore a more aggressive hyperparameter setting by doubling both learning rates and increasing the teacher score CFG scale to 6. These modifications yield substantial improvements in lip-sync accuracy at a modest cost to visual quality, with mouth movements becoming significantly more pronounced and articulated. However, we find that further increasing these parameters results in either training instability or visual oversaturation.

Final Configuration without Curated Multimodal Conditions. To verify the importance of curated conditions, we apply our other improvements to the baseline’s multimodal conditions. Results show that visual quality remains

degraded when distilling with lower-quality conditions, with persistent issues including poor temporal consistency and abrupt color shifts between frames, highlighting that data curation is essential for successful distillation.

Examples of generated videos from each configuration are shown in Figure 4. These results validate the effectiveness of each of our proposed improvement in enhancing generation quality and stabilizing multimodal distillation.

6 Conclusion

We propose an improved on-policy distillation recipe that distills a bidirectional, many-step video diffusion model into a causal, 4-step AR video diffusion. Our recipe incorporates refined multimodal conditioning, converged ODE initialization, and aggressive optimization to achieve strong visual fidelity and audio-visual synchronization while delivering sub-second first-frame latency and real-time throughput. Building on the distilled model, we develop LiveTalk, a real-time multimodal interactive avatar system that streams video conditioned on text, image, and audio. LiveTalk delivers coherent multi-turn interactions with significantly lower latency than state-of-the-art video models. These results enable new possibilities for real-time multimodal human-AI interactive systems.

References

- [1] Rishabh Agarwal, Nino Vieillard, Yongchao Zhou, Piotr Stanczyk, Sabela Ramos Garea, Matthieu Geist, and Olivier Bachem. 2024. On-policy distillation of language models: Learning from self-generated mistakes. In *The twelfth international conference on learning representations*.
- [2] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*.
- [3] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yanzhi Zhu, Mingkun Yang, Zhaoai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. 2025a. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- [4] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. 2025b. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- [5] David Berthelot, Arnaud Autef, Jierui Lin, Dian Ang Yap, Shuangfei Zhai, Siyuan Hu, Daniel Zheng, Walter Talbott, and Eric Gu. 2023. Tract: Denoising diffusion models with transitive closure time-distillation. *arXiv preprint arXiv:2303.04248*.
- [6] Tim Brooks, Bill Peebles, Connor Holmes, Will DePue, Yufei Guo, Li Jing, David Schnurr, Joe Taylor, Troy Luhman, Eric Luhman, et al. 2024. Video generation models as world simulators. *OpenAI Blog*, 1(8):1.
- [7] Jake Bruce, Michael D Dennis, Ashley Edwards, Jack Parker-Holder, Yuge Shi, Edward Hughes, Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, et al. 2024. Genie: Generative interactive environments. In *Forty-first International Conference on Machine Learning*.
- [8] Boyuan Chen, Diego Martí Monsó, Yilun Du, Max Simchowitz, Russ Tedrake, and Vincent Sitzmann. 2024. Diffusion forcing: Next-token prediction meets full-sequence diffusion. *Advances in Neural Information Processing Systems*, 37:24081–24125.
- [9] Brian Chen. 2025. [Question about image2video distillation implementation](#). GitHub issue. Issue #51, tianweiyl/-CausVid repository.
- [10] Yi Chen, Sen Liang, Zixiang Zhou, Ziyao Huang, Yifeng Ma, Junshu Tang, Qin Lin, Yuan Zhou, and Qinglin Lu. 2025. Hunyuanvideo-avatar: High-fidelity audio-driven human animation for multiple characters. *arXiv preprint arXiv:2505.20156*.
- [11] Joon Son Chung and Andrew Zisserman. 2017. [Out of time: Automated lip sync in the wild](#). pages 251–263.
- [12] Jiahao Cui, Hui Li, Yun Zhan, Hanlin Shang, Kaihui Cheng, Yuqi Ma, Shan Mu, Hang Zhou, Jingdong Wang, and Siyu Zhu. 2024. Hallo3: Highly dynamic and realistic portrait image animation with diffusion transformer networks. *arXiv e-prints*, pages arXiv–2412.
- [13] Ariel Ephrat, Inbar Mosseri, Oran Lang, Tali Dekel, Kevin Wilson, Avinatan Hassidim, William T Freeman, and Michael Rubinstein. 2018. Looking to listen at the cocktail party: A speaker-independent audio-visual model for speech separation. *arXiv preprint arXiv:1804.03619*.
- [14] Qijun Gan, Ruizhi Yang, Jianke Zhu, Shaofei Xue, and Steven Hoi. 2025. Omniavatar: Efficient audio-driven avatar video generation with adaptive body animation. *arXiv preprint arXiv:2506.18866*.
- [15] Zhifu Gao, Zerui Li, Jiaming Wang, Haoneng Luo, Xian Shi, Mengzhe Chen, Yabin Li, Lingyun Zuo, Zhihao Du, Zhangyu Xiao, and Shiliang Zhang. 2023. Funasr: A fundamental end-to-end speech recognition toolkit. In *INTERSPEECH*.
- [16] Xingtong Ge, Xin Zhang, Tongda Xu, Yi Zhang, Xinjie Zhang, Yan Wang, and Jun Zhang. 2025. [Senseflow: Scaling distribution matching for flow-based text-to-image distillation](#).
- [17] Google DeepMind. 2025. [Veo 3 technical report](#).
- [18] Jiatao Gu, Shuangfei Zhai, Yizhe Zhang, Lingjie Liu, and Joshua M Susskind. 2023. Boot: Data-free distillation of denoising diffusion models with bootstrapping. In *ICML 2023 Workshop on Structured Probabilistic Inference* {&} *Generative Modeling*, volume 3.
- [19] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. 2018. [Gans trained by a two time-scale update rule converge to a local nash equilibrium](#).

- [20] Jonathan Ho and Tim Salimans. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.
- [21] Teng Hu, Zhentao Yu, Zhengguang Zhou, Sen Liang, Yuan Zhou, Qin Lin, and Qinglin Lu. 2025. Hunyuancustom: A multimodal-driven architecture for customized video generation. *arXiv preprint arXiv:2505.04512*.
- [22] Xun Huang, Zhengqi Li, Guande He, Mingyuan Zhou, and Eli Shechtman. 2025. Self forcing: Bridging the train-test gap in autoregressive video diffusion. *arXiv preprint arXiv:2506.08009*.
- [23] Xuan Ju, Weicai Ye, Quande Liu, Qiulin Wang, Xintao Wang, Pengfei Wan, Di Zhang, Kun Gai, and Qiang Xu. 2025. Fulldit: Multi-task video generative foundation model with full attention. *arXiv preprint arXiv:2503.19907*.
- [24] Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- [25] Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, et al. 2024. Hunyuanvideo: A systematic framework for large video generative models. *arXiv preprint arXiv:2412.03603*.
- [26] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th symposium on operating systems principles*, pages 611–626.
- [27] Seongyun Lee, Seungone Kim, Sue Park, Geewook Kim, and Minjoon Seo. 2024. Prometheus-vision: Vision-language model as a judge for fine-grained evaluation. In *Findings of the association for computational linguistics ACL 2024*, pages 11286–11315.
- [28] Mingxiao Li, Bo Wan, Marie-Francine Moens, and Tinne Tuytelaars. 2024. Animate your motion: Turning still images into dynamic videos. In *European Conference on Computer Vision*, pages 409–425. Springer.
- [29] Zongyu Lin, Wei Liu, Chen Chen, Jiasen Lu, Wenzhe Hu, Tsu-Jui Fu, Jesse Allardice, Zhengfeng Lai, Liangchen Song, Bowen Zhang, et al. 2025. Stiv: Scalable text and image conditioned video generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16249–16259.
- [30] Kunhao Liu, Wenbo Hu, Jiale Xu, Ying Shan, and Shijian Lu. 2025. Rolling forcing: Autoregressive long video diffusion in real time. *arXiv preprint arXiv:2509.25161*.
- [31] Chetwin Low and Weimin Wang. 2025. Talkingmachines: Real-time audio-driven facetime-style video via autoregressive diffusion models. *arXiv preprint arXiv:2506.03099*.
- [32] Kevin Lu and Thinking Machines Lab. 2025. [On-policy distillation](https://thinkingmachines.ai/blog/on-policy-distillation). *Thinking Machines Lab: Connectionism*. <https://thinkingmachines.ai/blog/on-policy-distillation>.
- [33] Mang Ning, Mingxiao Li, Jianlin Su, Albert Ali Salah, and Itir Onal Ertugrul. 2023. Elucidating the exposure bias in diffusion models. *arXiv preprint arXiv:2308.15321*.
- [34] OpenAI. 2025. [Sora 2 is here](#). Research blog.
- [35] William Peebles and Saining Xie. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4195–4205.
- [36] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695.
- [37] Florian Schmidt. 2019. Generalization in generation: A closer look at exposure bias. *arXiv preprint arXiv:1910.00292*.
- [38] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. 2023. Consistency models.
- [39] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.
- [40] Hansi Teng, Hongyu Jia, Lei Sun, Lingzhi Li, Maolin Li, Mingqiu Tang, Shuai Han, Tianning Zhang, WQ Zhang, Weifeng Luo, et al. 2025. Magi-1: Autoregressive video generation at scale. *arXiv preprint arXiv:2505.13211*.
- [41] Thomas Unterthiner, Sjoerd van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. 2019. [Towards accurate generative models of video: A new metric & challenges](#).

- [42] GitHub User. 2024. [Question about ode_init.pt weight and dmd training](#). GitHub Issue. Issue #50, Self-Forcing repository.
- [43] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, et al. 2025. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*.
- [44] Mengchao Wang, Qiang Wang, Fan Jiang, Yaqi Fan, Yunpeng Zhang, Yonggang Qi, Kun Zhao, and Mu Xu. 2025. Fantasytalking: Realistic talking portrait generation via coherent motion synthesis. In *Proceedings of the 33rd ACM International Conference on Multimedia*, pages 9891–9900.
- [45] Huawei Wei, Zejun Yang, and Zhisheng Wang. 2024. Aniportrait: Audio-driven synthesis of photorealistic portrait animation. *arXiv preprint arXiv:2403.17694*.
- [46] Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, Kun Yan, Sheng-ming Yin, Shuai Bai, Xiao Xu, Yilei Chen, et al. 2025. Qwen-image technical report. *arXiv preprint arXiv:2508.02324*.
- [47] Haoning Wu, Zicheng Zhang, Weixia Zhang, Chaofeng Chen, Liang Liao, Chunyi Li, Yixuan Gao, Annan Wang, Erli Zhang, Wenxiu Sun, Qiong Yan, Xiongkuo Min, Guangtao Zhai, and Weisi Lin. 2023. [Q-align: Teaching lmms for visual scoring via discrete text-defined levels](#).
- [48] Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. 2024. [Efficient streaming language models with attention sinks](#).
- [49] Jin Xu, Zhifang Guo, Hangrui Hu, Yunfei Chu, Xiong Wang, Jinzheng He, Yuxuan Wang, Xian Shi, Ting He, Xinfu Zhu, et al. 2025. Qwen3-omni technical report. *arXiv preprint arXiv:2509.17765*.
- [50] Shuai Yang, Wei Huang, Ruihang Chu, Yicheng Xiao, Yuyang Zhao, Xianbang Wang, Muyang Li, Enze Xie, Yingcong Chen, Yao Lu, et al. 2025. Longlive: Real-time interactive long video generation. *arXiv preprint arXiv:2509.22622*.
- [51] Tianwei Yin, Michaël Gharbi, Taesung Park, Richard Zhang, Eli Shechtman, Fredo Durand, and Bill Freeman. 2024a. Improved distribution matching distillation for fast image synthesis. *Advances in neural information processing systems*, 37:47455–47487.
- [52] Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Fredo Durand, William T Freeman, and Taesung Park. 2024b. One-step diffusion with distribution matching distillation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6613–6623.
- [53] Tianwei Yin, Qiang Zhang, Richard Zhang, William T Freeman, Fredo Durand, Eli Shechtman, and Xun Huang. 2025. From slow bidirectional to fast autoregressive video diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 22963–22974.
- [54] Zhimeng Zhang, Lincheng Li, Yu Ding, and Changjie Fan. 2021. Flow-guided one-shot talking face generation with a high-resolution audio-visual dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3661–3670.
- [55] Shangchen Zhou, Kelvin C. K. Chan, Chongyi Li, and Chen Change Loy. 2022. [Towards robust blind face restoration with codebook lookup transformer](#).
- [56] Hao Zhu, Wayne Wu, Wentao Zhu, Liming Jiang, Siwei Tang, Li Zhang, Ziwei Liu, and Chen Change Loy. 2022. Celebv-hq: A large-scale video facial attributes dataset. In *European conference on computer vision*, pages 650–667. Springer.

A Multi-Round Evaluation Details

This appendix provides a comprehensive description of the nine evaluation dimensions comprising our interaction benchmark. Each dimension is assessed by a VLM using carefully designed, structured prompts that define specific scoring criteria. The evaluation framework is divided into two major categories: **Visual Interaction Performance** and **Interaction Content Quality**. We also provide extra implementation details for the evaluation protocol.

A.1 Visual Interaction Performance

Visual interaction performance captures the non-verbal and emotional aspects of the assistant’s video responses, emphasizing the quality of human-like engagement and visual consistency across conversational turns.

Emotional Appropriateness. This dimension evaluates the degree to which the assistant’s facial expressions align with conversational context and emotional content.

Assessment criteria include:

- **Content-Emotion Matching:** Appropriateness of emotional expressions relative to semantic content (e.g., conveying seriousness during complex explanations; displaying warmth during positive encouragement).
- **Intensity Calibration:** Proportionality of emotional intensity to the significance and tone of the conversational content.
- **Emotional Transitions:** Naturalness and coherence of emotional shifts across multiple consecutive video responses.
- **User Responsiveness:** Adaptive emotional responses that acknowledge and appropriately react to the user’s inferred emotional state (e.g., providing reassurance for confusion, matching enthusiasm).

Nonverbal Interaction. This dimension quantifies the quality of interactive nonverbal cues that establish and maintain user engagement.

Assessment criteria include:

- **Eye Contact Quality:** Maintenance of natural, engaging gaze patterns that simulate direct interaction with the user.
- **Interactive Gestures & Micro-expressions:** Effective use of confirmatory nods, thoughtful pauses, raised eyebrows, and other subtle movements that signal active participation in dialogue.
- **Listening Response:** Presence of nonverbal acknowledgment signals that indicate comprehension and processing of user input prior to verbal response.
- **Interaction Authenticity:** Overall impression of genuine bidirectional dialogue rather than unidirectional content delivery.

Multi-Video Coherence. This dimension assesses visual and behavioral consistency across multiple video responses within the conversation. This dimension is important as it evaluates the immersion and believability in multi-round interactions, which is especially important in extended interactions.

Assessment criteria include:

- **Visual Identity Consistency:** Stability of the speaker’s physical appearance, including facial features, hairstyle, and attire.
- **Scene Consistency:** Stability of environmental factors including background elements, lighting conditions, and camera perspective.
- **Emotional Continuity:** Logical progression and smooth evolution of emotional states from one video response to subsequent responses.
- **Behavioral Coherence:** Consistency in posture, gestural patterns, speech cadence, and overall presentation style.

Conversational Naturalness. This dimension measures the degree to which the interaction conveys authentic human-like spontaneity rather than scripted or mechanistic behavior.

Assessment criteria include:

- **Dialogue vs. Broadcast Distinction:** Presence of interactive, bidirectional conversational dynamics versus formal, unidirectional presentation style.
- **Spontaneous Human Characteristics:** Natural occurrences of brief pauses, thinking moments, speech hesitations, and other markers of authentic spontaneous communication.

A.2 Interaction Content Quality

Interaction content quality evaluates the semantic, logical, and contextual appropriateness of the assistant’s verbal responses, ensuring both accuracy and conversational coherence.

Semantic Relevance. This dimension assesses whether the assistant’s responses directly and accurately address the topical focus of user queries.

Information Completeness. This dimension evaluates if the provided answers are comprehensive, sufficiently detailed, and self-contained.

Logical Consistency. This dimension identifies contradictions, factual errors, or logical inconsistencies both within individual responses and across the entire conversational history.

Context Understanding. This dimension measures the assistant’s ability to maintain conversational state, correctly resolve anaphoric references, and effectively build upon information established in prior dialogue turns.

Overall Interaction Experience. This dimension provides a holistic assessment encompassing conversational fluency, effectiveness in task completion, and overall user satisfaction.

A.3 Evaluation Protocol Details

VLM Evaluator. We employ Qwen2.5-VL-30B-A3B-Instruct (Bai et al., 2025a) as our evaluation model, deployed via vLLM (Kwon et al., 2023). The model configuration supports contexts up to 256K tokens and enhanced video understanding capabilities, processing up to 32 frames per video at 784×784 resolution. To ensure stable and consistent scoring, we use a low sampling temperature of 0.1 combined with top-p sampling ($p=0.9$). Each evaluation dimension is guided by a detailed rubric encoded in dimension-specific prompt files, with model responses structured as `\boxed{score}` for automated parsing and aggregation.

ASR Transcription. Full conversational transcripts are constructed by transcribing both user audio inputs (16kHz WAV/MP3 format) and assistant video audio tracks into English text. We utilize the FunASR (`speech-paraformer-asr-en-16k-vocab4199`) (Gao et al., 2023) model, which employs a non-autoregressive Paraformer architecture optimized for high-accuracy real-time transcription.

Multi-Video Evaluation. For conversations containing multiple assistant video responses, all video files are presented to the VLM evaluator in chronological order alongside the complete conversational transcript. This multi-video presentation enables both isolated analysis of individual video segments and comparative cross-video analysis, which is essential for accurately assessing the *Multi-Video Coherence* dimension and evaluating *Emotional Transitions*. All content quality dimensions leverage the full conversation history to assess semantic coherence, logical consistency, and contextual understanding across all dialogue turns.

B Training Configuration Details

The distillation pipeline consists of two sequential stages: ODE initialization and on-policy distribution matching distillation (DMD) (Section 3.1). Tab. 4 details the hyperparameters and optimization settings for both stages.

Table 4: **Training configurations for ODE trajectory initialization and DMD distillation.** CFG (Generating Trajectory) denotes the classifier-free guidance scale used by the bidirectional model (OmniAvatar-1.3B) during ODE trajectory generation. For DMD distillation, the critic score network s_ψ is first updated for 20 steps to obtain accurate score estimates, after which the student generator network g_ϕ is trained (Section 3.1). Exponential moving average (EMA) updates are enabled from step 200 onward.

Hyperparameters	ODE Initialization
Batch size	64
Learning rate	$4e-5$
Optimizer	AdamW, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1e-8$, weight_decay=0
CFG (Generating Trajectory)	4.5
Hyperparameters	DMD Distillation
Teacher score network	OmniAvatar-14B
Teacher score CFG	6.0
Critic score network	OmniAvatar-1.3B
Batch size	64
Optimizer (g_ϕ)	AdamW, $\beta_1 = 0$, $\beta_2 = 0.999$, $\epsilon = 1e-8$, weight_decay=0.01
Optimizer (s_ψ)	AdamW, $\beta_1 = 0$, $\beta_2 = 0.999$, $\epsilon = 1e-8$, weight_decay=0.01
Learning rate (g_ϕ)	$4e-6$
Learning rate (s_ψ)	$8e-7$
Generator/critic update ratio	5, with critic warmup for 20 steps
EMA decay	0.99